

VIRTEX-5QV ARCHITECTURAL FEATURES SEU CHARACTERIZATION SUMMARY

Gary Swift and Carl Carmichael
Xilinx, Inc.
San Jose, California

Gregory Allen
Jet Propulsion Laboratory
Pasadena, California

Roberto Monreal
Southwest Research Institute
San Antonio, Texas

George Madias and Eric Miller
Boeing
El Segundo, California

with help from other members of the
Xilinx Radiation Test Consortium

Document Revision History

Initial Release: April 30, 2013

First Revision: August 23, 2013

General cleanup and “TBD” eradication. Chapter 8 with calculated space rates for example orbits added. Proton data on User Flip-flops added to Section 3.2 and heavy ion data fit parameters for IOSERDES added into Section 6.3. Elaborated on upset signatures for DCM and PLL tests in Chapter 7 and IODELAY test and results in Section 6.4.

A portion of the work reported here was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the National Aeronautics and Space Administration Electronic Parts and Packaging (NEPP) Program.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

Copyright 2012-2013. All rights reserved.

TABLE OF CONTENTS

1	OVERVIEW	1
2	BRAM ECC	2
3	USER FLIP-FLOPS AND SET FILTERS	4
	3.1 Overview	4
	3.2 Test Results	4
4	DIGITAL SIGNAL PROCESSING (DSP) BLOCKS	7
	4.1 Overview	7
	4.2 Dynamic Test Results	7
	4.3 Dynamic-Static Comparison over Frequency	9
5	MULTIGIGABIT TRANCEIVER (MGT) BLOCKS.....	10
	5.1 Overview	10
	5.2 Raw MGT Test Results.....	11
	5.3 Aurora Protocol Test Results	12
	5.4 RapidIO Protocol Test Results.....	12
	5.5 Comparing Raw and Protocol Test Results	13
6	INPUT/OUTPUT (I/O)-RELATED BLOCKS.....	15
	6.1 LVCMOS I/Os	15
	6.2 LVDS I/Os	17
	6.3 IOSERDES	24
	6.4 IODELAY	26
7	CLOCK MANAGEMENT (CMT).....	29
	7.1 Digital Clock Management (DCM) Blocks	32
	7.2 Phase-Locked Loop (PLL) Blocks.....	35
8	SPACE UPSET RATES FOR SELECTED ORBITS	38
9	REFERENCES	42

1 OVERVIEW

This report summarizes Xilinx Radiation Test Consortium (XRTC) Single-Event Effects (SEE) testing of the space-grade XQR5VFX130 FPGA's main architectural function blocks. It is a companion to the Static Summary Report [1] on the XRTC's static SEU tests of the same device, the Virtex-5QV. The tested architectural features are powerful and complex, offering a lot of programmable features and options. As a result, the tests and data analysis can be quite complex. Of necessity, the test matrices cannot provide full coverage of all possible operating modes. This Architectural Features Summary report only describes at a high level (a) the feature or block of interest, (b) the test methodology employed and (c) the most important results obtained in order to enable calculation of expected and/or worst-case SEE rates on those blocks given an arbitrary on-orbit radiation environment. It is also hoped that identifying and quantifying the main residual error modes will inspire mitigation ideas for the Consortium to test.

More complete documentation on the blocks tested is available, usually in the form of a User's Guide on the Xilinx website, www.xilinx.com. Detailed test reports are available for some of these architectural features and they are highlighted in each feature's "Overview" chapter and collected in the "References." JPL has collected many of these on its FPGA website <http://parts.jpl.nasa.gov/organization/group-5144/radiation-effects-in-fpgas/xilinx/>.

The Virtex-5QV uses radiation-hard-by-design (RHBD) techniques to lower Single-Event Upset (SEU) susceptibility on key memory elements. In particular, the dual node configuration cells require a minimum charge collection at two intentionally widely spaced nodes; they achieve about 1000x improvement over the earlier space-grade Virtex-4QV [2]. This eliminates the main source of errors previously seen in the various architectural blocks [3] revealing a new underlying layer of more fundamental error signatures. In addition the RHBD techniques were applied to significantly reduce Single-Event Functional Interrupts (SEFIs) achieving results, documented in Ref. 1, that are about 100x lower than the Virtex-4QV in geosynchronous orbit (GEO).

This testing Summary is leveraged from the combined efforts of the members of the XRTC, occasionally and, perhaps, more accurately known as the Xilinx SEE Test Consortium. The XRTC is a voluntary association of aerospace entities, including leading aerospace companies, universities, space agencies and national laboratories, combining resources to characterize reconfigurable FPGAs for aerospace applications. Previous presentations and publications of Virtex-5QV radiation results have been made by Consortium members, notably at the NSREC and MAPLD conferences and the SEE Symposium, as well as at our own XRTC Annual Meeting. A parallel report for the Virtex-4QV family devices is the Dynamic and Mitigation Test Report [3].

2 BRAMECC

An enhanced Error Correcting Code (ECC) block is included in Virtex-5 family devices [4a] in conjunction with the provided block RAM (BRAM). The Virtex-5QV ECC has added an optional automatic write-back feature so that upsets are corrected even when only reading the (corrected) data. Hamming-code-based ECC is well known to be an effective way to prevent system errors in the presence of a few random upsets. The write-back feature makes it easier to prevent upset accumulation which would eventually overwhelm the ECC. The Virtex-5QV ECC uses sets of 8 check bits to protect 64-bit data words; the performance penalty for the extra robustness it provides is documented in the Virtex-5QV Electrical Characteristics Data Sheet, DS692 [4b].

The XRTC subjected the Virtex-5QV to beam testing to establish the ECC effectiveness [5]. However, beam testing creates upsets much faster than a real space radiation environment, even a severe one; thus extrapolation down to expected upset rates is necessary. Guided by the known statistical relationship between the underlying upset rate and the system error rate, the data plotted in Figure 1 allow such an extrapolation. At high beam fluxes, the data follow a quadratic relationship showing effective mitigation, but as the upset rate is lowered a transition to a linear relationship with system errors is seen indicating that single points of failure dominate the extrapolation. Presumably, those points reside in the ECC block in the form of buffer upsets and logic transients. However, the number of failure points is small enough that the projected error rates on-orbit are very low, likely to meet even the most severe requirements.




Figure 1. BRAM ECC Test Results as a Function of Upset Rate [Ref. 5, Fig.3]

This mitigation test was conducted over many orders of magnitude of flux with normal incidence ions to measure the single-points-of-failure in the ECC block because that will

dominate the errors seen at space irradiation rates which are many orders of magnitude off the scale of Figure 1 to the left. There is, however, a legitimate concern that a single ion causing multiple bits to upset (MBUs) would also defeat the error correction capabilities. Ions are more likely to upset multiple cells if they strike at a large angle to the silicon face, but aligned with rows or columns of the memory array. However, bits in the same ECC word are not adjacent which alleviates the MBU concern. This result is buttressed by the high angle BRAM experiment documented in the Virtex-5QV static report [1].

In summary, there are three possible ways to get BRAM errors when the ECC feature is used: 1) two (or more) coincident upset events that appear in a single 72-bit correction word within a single scrub cycle, 2) an individual ion strike that upsets two (or more) bits in a 72-bit word, or 3) upsets or transients in the ECC circuitry itself while a write or read is occurring. For any reasonable scrub time, the chance of coincident separate upsets is negligible while the physical layout of the bits prevents MBUs from having any significant chance of occurring within a correction word. Thus, single-points-of-failure in the ECC circuitry will dominate the observed error rate in space radiation environments and the solid curve of Figure 1 can be used to extrapolate to a given environment using a slope of 1. For example, in geosynchronous orbit (GEO), the BRAM upset rate for all the bits is about 3/day [1]. That is roughly 3×10^{-12} upsets per bit-sec or about five and a half orders of magnitude lower than the left end of the x-axis in Figure 1. Down from the y-intercept by that amount yields 1×10^{-9} ECC errors per second or an average interval between errors of more than 30 years. This shows that the single-points-of-failure in the ECC circuitry are relatively few in number.

3 USER FLIP-FLOPS AND SET FILTERS

The Virtex-5 FX130 FPGAs provide more than 80kb of user registers [4c]. In the space-grade XQR5VFX130, these are implemented with a radiation-hard-by-design (RHDB) approach where both the master and slave flip-flops have “geometric” upset hardening similar to what was used for configuration cells. Each flip-flop half of a user register bit has dual separated nodes that must both collect charge in order for that flip-flop to settle in an upset state. Thus, the direct upsetability of the bits is very low, to the point where the indirect mechanism of single-event transients on the inputs is clearly dominant.

In anticipation of that, single-event transient filters on all flip-flop inputs, that is, all data, clock, and control lines, are available. These filters are implemented via duplicate, but time-shifted inputs so that only transients greater than about 800ns will be seen by the flip-flop inputs. Because the extra upset hardness the filters provide does come with a performance penalty, the use of the SET filters is at the space designer’s option either via a global selection or by individually selecting tiles.

3.1 Overview

In-beam XRTC testing of the flip-flops incorporated a fairly extensive matrix of conditions, resulting in a large number of test results. First, the obvious dependences were investigated using variations on a shift register design that filled 90+% of the device: SET filters on and off, data line frequency (1.5, 100 and 200 MHz) and data pattern (all ones, all zeros, and checkerboard). In addition, the XRTC test matrix included three different amounts of intervening logic (zero, one or four levels) and the notion of extra LUT sensitivity (or not) depending on the exact implementation of “don’t care” logic associated with unused LUT inputs. Data sets were collected on all one hundred and eight ($=2 \times 3 \times 3 \times 3 \times 2$) separate combinations. Some of the heavy ion test results were previously presented at the SEE Symposium [6, 7]. The test design from George Madias of Boeing incorporates 24 separate shift register chains of 800 flip-flops each. The chains are individually monitored so that simultaneous data is collected on six or eight of the combinations, allowing cleaner comparisons without inherent run-to-run dosimetry variations.

3.2 Test Results

The test results are quite clear. First, the effectiveness of the SET filters is evident; see Fig. 2 where there is about an order of magnitude in susceptibility difference at all LETs except near the threshold. Next, the expected linear frequency dependence is clearly demonstrated when the worse-case pattern of checkerboard is used; this is attributed to clock tree SETs. For intervening logic, we conclude that flip-flop and logic output drivers are bigger sources of SETs than inputs. A small filtering effect from routing would explain why adding targets by adding intervening logic in parallel (adding width) is worse than stacking intervening logic serially. The all zeros pattern is least susceptible being sensitive only to direct upset (which are very low due to RHBD) plus data line hits that are in coincidence with the clock edge. More susceptible than all zeros, the all ones pattern adds sensitivity to asynchronous hits on the reset inputs. The checkerboard pattern is the most upset-able pattern; it adds clock tree SETs which more than compensate for the 50% reduction in upsets caused by reset SETs. Finally and somewhat surprisingly, “sensitive” LUTs were only very slightly more sensitive; this is because although they

increase the number of critical bits in the configuration significantly those bits are upset hard by design and, thus, don't contribute much to the error rate.




Figure 2a. Filter Effectiveness Comparison: Heavy-ion F/F Results as a Function of LET

As indicated in the legend of Fig. 2a with units of upsets per bit-day, CREME96 [8] was used by Boeing to calculate flip-flop upset rates for the case of geosynchronous orbit. Note that these numbers were generated with CREME parameters that deviate from the XRTC norm in a couple of minor areas relating to the RPP volume: x and y dimensions are set as the square root of the cross section at an LET of 80 and the sensitive depth is set to 2 microns with a half micron funnel depth. GEO rates calculated with the usual XRTC CREME settings are given in Chapter 8. Either way, it is clear that the RHBD upset hardening of the flip-flops is quite effective, about 1000x in comparison to the unhardened registers of the DSPs (next Chapter) or to the BRAMs (previous Chapter and Section 3.3.1 of the Static Report) for example. Turning on the all-input SET filters further improves the susceptibility by more than an order of magnitude.

Table 1. Weibull Fit Parameters for Virtex-5QV User Flip-Flop Upsets from Heavy Ions

	Weibull Parameters			
	Limit (cm²/bit)	Onset (MeV-cm²/mg)	Width -	Power -
worse-case				
filters=OFF	9.78x10 ⁻⁸	0.67	26	1.76
filters=ON	3.45x10 ⁻⁸	0.25	86	1.75




Figure 2b. Filter Effectiveness Comparison: Proton F/F Results as a Function of Energy

The proton response of the upset-hardened flip-flops is quite small even with the SET filters off. Flip-flop test data was taken on the Jan and June 2012 test dates in the Nov 2011 to Dec 2012 proton campaign and concentrated data coverage on the previously identified worst-case frequency and pattern conditions of the test matrix- 200 MHz and checkerboard- with filters both on and off for three combinations of intervening logic types- none and parallel using four LUT (look-up table logic element) invertors in maximum and minimum sensitivity constructions. While some events affect multiple chains, results shown are based on the more conservative approach of using the full flip-flop upset counts, not the somewhat reduced counts of events.

Table 2. Weibull Fit Parameters for Virtex-5QV User Flip-Flop Upsets from Protons

worse-case (200 MHz, pattern:checkerboard)	Weibull Parameters			
	Limit (cm²/bit)	Onset (MeV)	Width -	Power -
filters=OFF	4.0×10^{-15}	8	30	1
filters=ON	2.7×10^{-16}	8	30	1

Example on-orbit rates were calculated with CREME96 [8] and are presented in Chapter 8 . For most orbits, the RHBD flip-flops experience a lower upset rate than in GEO as the reduction in heavy ions from geomagnetic shielding more than compensates for the addition of trapped protons, especially for the case where the SET filters are ON.

4 DIGITAL SIGNAL PROCESSING (DSP) BLOCKS

The Digital Signal Processing (DSP) Slices provide advanced high-speed arithmetic and comparison functions, including multiply and accumulate. These may be strung together to rapidly perform more complex calculations, like Fast Fourier Transforms, on continuous streams of data. For details on the DSP blocks, including basic function, speed, available ports and op codes, consult the DSP User's Guide, UG193 [9].

4.1 Overview

The purpose of the dynamic testing described here is to characterize the Virtex-5QV DSP blocks in operation and over frequency. Virtex-5 DSPs have forty-two defined operations or op codes; only a representative subset consisting of three op codes were actually used in these tests, addition, multiplication, and accumulation, with a fixed set of inputs. In addition the DSP blocks incorporate several unhardened optional input, output and pipeline registers that will contribute to the error rates observed in these dynamic tests. Recalling that the static upset characteristics of the registers are documented in Ref [1], it is possible to account for the registers and duty cycles used and calculate the static component in the dynamic tests. Beyond the high level summary presented below, full details of this calculation as well as the dynamic test methodology and results have previously been published [10] and recently re-visited and extended [11] and now includes data taken at the Texas A&M cyclotron in Sept. 2013.

4.2 Dynamic Test Results

Errors observed in the multiplication and addition tests are most commonly of short duration since new operands are loaded after each cycle. However, errors in the accumulation test tend to be persistent because earlier errors appear in the current operands. Two of the measured susceptibility curves and their fits are shown in Figure 3 and all three fits are parameterized in Table 3 for heavy ions. More details, including breaking the total errors into the various error signatures observed is included in the full DSP testing report [10] that this chapter and its static report [1] counterpart summarize.




Figure 3a. Heavy Ion Results for DSPs Performing Accumulation [Ref. 10, Fig. 7.11 + more data]




Figure 3b. Heavy Ion Results for DSPs Performing Multiplication [Ref. 10, Fig. 7.14 + more data]

Heavy ion data was taken at both Berkeley's 88" Cyclotron and at Texas A&M Cyclotron Institute using three frequencies: 6.25, 12.5, and 25 MHz, all included in Figure 3. The 25 MHz Berkeley data is clearly inconsistent with the rest of the data; see the two low outlier points in Figure 3b for example. Additional low LET data was collected at Texas A&M in Sept. 2012 and the accumulate op code fit in Table 3 is improved over what's reported in Ref 10.

Table 3. Weibull Fit Parameters for Virtex-5QV DSP Upsets from Heavy Ions

op-code	Weibull Parameters			
	Limit (cm ² /DSP)	Onset (MeV-cm ² /mg)	Width -	Power -
multiply	3.65x10 ⁻⁵	0.1	89	1.11
addition	1.35x10 ⁻⁵	0.1	44	1.08
accumulate	3.28x10 ⁻⁶	0.4	15	1.18

Proton data was collected with the DSPs operating at 6.25 MHz in December 2012 at the UC-Davis cyclotron for two energies; that data is shown in Figure 4. Note that we lack data at low energy to define the threshold better. Further, the Consortium has not undertaken the effort to push the frequency higher in order to experimentally characterize the frequency dependence.




Figure 4. Proton Results for DSPs for Three Op Codes Run at 6.25 MHz

For upset-hardened elements like the configuration bits, GEO is the worst-case orbit; for lower orbits, the reduction of heavy ion effects due to increased geomagnetic shielding will be more than compensated for by the addition of trapped protons. However, that is

not the case for unhardened features like the DSP registers and proton upsets will dominate the error rates for proton-rich orbits; see the examples in Chapter 8.

Table 4. Weibull Fit Parameters for Virtex-5QV DSP Upsets from Protons

op-code	Weibull Parameters			
	Limit (cm ² /DSP)	Onset (MeV)	Width -	Power -
multiply	1.10x10 ⁻¹¹	4.0	10.2	0.298
addition	8.0x10 ⁻¹²	4.0	19.4	0.901
accumulate	4.0x10 ⁻¹²	4.0	37	0.80

4.3 Dynamic-Static Comparison over Frequency

Dynamic test results include a frequency-independent component (measured separately in the static tests [1]) as well as a frequency-dependent component. Subtracting out the static component from the dynamic data presented above yields the frequency-dependent component at about 25 MHz. Table 5 presents GEO rate calculations separating out the frequency-independent (or static) and frequency-dependent (or dynamic) components for the three op codes tested.

Table 5. Relative Static & Dynamic Contributions to Virtex-5QV DSP Upset Rate

Orbit: GEO op code	Static Rate* (upsets/dev-yr)	Dynamic Rate*t (upsets/dev-yr)	Static Fraction-
multiply	11.8	21.3	55%
addition	5.90	18.4	32%
accumulate	2.02	4.42	45%

* - Assumes all 320 DSP blocks per device are used and 25 MHz

Note that for this analysis, the CREME parameters deviate slightly from the XRTC norm: shielding of 150 mils of aluminum (rather than 100) and choosing to include Z's of 2 to 92 (instead of 1-92).

To extrapolate to higher frequency, a linear model is reasonable where the dynamic fraction is scaled with frequency and the static fraction is added. For example, because the multiplication result is about half frequency independent, extrapolating to 250 MHz (or a factor of 10x higher frequency) will result in an error rate that is expected to be about 5.5x higher than observed at 25 MHz (ten times the half that scales plus the half that doesn't). In other words, a DSP block performing multiplication would be erroneous due to an upset in GEO about four times a decade, on average.

5 MULTIGIGABIT TRANCEIVER (MGT) BLOCKS

High-speed serial receiver/transmitter pairs are finding more and more use in ground-based electronics. Marking the first offering in a space-grade FPGA, there are eighteen high speed (up to 4.25 Gb/s) serial transceivers available in the Virtex-5QV. Xilinx designates these as multi-gigabit transceivers or MGTs and offers the higher performance type, dubbed GTX, in the Virtex-5QV and its commercial counterpart the XC5VFX130T. Each GTX transceiver supports full-duplex, with built-in circuits for clock-and-data recovery; a User Guide is available that provides full details of their capabilities and how to use them [12]

5.1 Overview

High-speed serial links typically run a “protocol” which defines how the information they convey is packetized and adds packet integrity checks and ability to request a bad packet be retransmitted. There are many choices of protocol with different advantages and disadvantages, like RapidIO or Ethernet, and it is impossible to test them all. Instead, the XRTC testing was first aimed at how heavy ions affect the underlying bit communication in the absence of any protocol. This fundamental susceptibility information can then be mapped onto any particular protocol to predict how radiation environments will affect bit error rates and availability for the combination of GTXs + the given protocol. In order to demonstrate that this prediction from fundamental data works, two selected protocols were also beam tested- RapidIO and a more straight-forward Xilinx-defined one, dubbed Aurora.

The sans-protocol (or “raw”) testing was done with a bit rate of 3.25 Gb/s; 10-bit symbols encode 8 bits of data, yielding an actual data rate of 2.6 Gb/s . The testing used an innovative spectroscopy-like technique invented by Roberto Monreal of the Southwest Research Institute where the time a particular error lasted was measured and, thus, event durations were available to help classify upset event types. A full report on the raw transceiver testing is available online [13] and a good summary has been published [14]

The RapidIO testing and the Aurora testing also used a 3.25 Gb/s clock rate. The addition of a protocol would generally be expected to add robustness, that is, increase the availability and/or decrease the effective bit error rate (BER) and the in-beam testing confirms that expectation. In addition, it is possible to improve robustness more with some extra effort to make the protocols more resistant to SEE. Improvements tested include incorporating the Virtex-5QV’s improved BRAM-ECC hard block into the buffers used by the RapidIO IP and, in the Aurora IP, implementing a supervisor function that detects the major error modes and resets them. Beam testing the latter shows that BER is somewhat improved and availability is significantly improved while the event rate is, of course, the same; these results are so successful in making the MGTs plus mitigated protocol robust that the remaining intrusive event rates are comparable to the extremely low SEFI rates (once in ten thousand years in GEO), indicating that further improvement efforts are not worthwhile.

5.2 Raw MGT Test Results

While the XRTC testing kept track of a large number of upset and recovery signatures, they fall into two main error baskets: 1) bit errors (BEs) and 2) loss-of-link events (LOLs). Bit errors can be from hits on either the transmitter or the receiver and the test methodology is able to clearly distinguish these. As can be seen in Figure 5a, the susceptibilities are very nearly the same with receiver errors being consistently a little more likely for all LETs. A single erroneous bit is by far the most common type of error seen under irradiation but, offsetting that, it is the least impactful in that it does not interrupt operation. On the other hand, loss-of-link is more serious in terms of the effect on bit error rate and availability and requires outside action, like resetting the receiver or transmitter or even the whole tile.




Figure 5a. 3.25 GHz MGT bit errors [14]




Figure 5b. 3.25 GHz MGT TX loss-of-link (LOL) errors [14]

The complete report [13] covers full details of the test methodology and the results observed, including time and recovery signatures as well as more detailed bit-error-rate (BER) and availability estimates. The results in Tables 6 simplifies and condenses the heavy ion results into fewer buckets. It's not necessarily clear from the fit parameters, but the rate of Rx Loss-of-link ends up being higher than Tx Loss-of-Link; similarly, Rx bit errors are somewhat more likely than Tx bit errors.

Table 6. Weibull Fit Parameters for Virtex-5QV MGT Upsets from Heavy Ions

error type	Weibull Parameters			
	Limit (cm²/MGT)	Onset (MeV-cm²/mg)	Width -	Power -
Rx Bit Errors	7.42×10^{-5}	0.01	36	1.39
Tx Loss-of-Link	1.57×10^{-5}	0.01	43	1.56
Rx Loss-of-Link	3.90×10^{-5}	0.01	66	1.42

5.3 Aurora Protocol Test Results

A full report is available [15] and a summary was recently presented [16]; this testing was conducted by Consortium member Brigham Young University originally in heavy ions. The results show that a protocol can add to the robustness inherent in the MGTs. Additionally, it was recognized that a small supervisor might be added on to enhance the robustness further. Both proton and heavy ion irradiations have aided the development and proven the effectiveness of this enhancement.

5.4 RapidIO Protocol Test Results

A second MGT-based protocol was also tested, the Xilinx implementation of the industry standard RapidIO protocol. The raw MGT results are quite good with years expected between errors and many decades between loss-of-link errors. Adding a protocol helps lower bit errors and increases availability, as was seen with the Aurora results. This is also true for the Xilinx RapidIO IP whose results are shown in Figure 6 as the uppermost “curve” with black squares as the data markers.

In addition Figure 6 shows some mitigation investigation results. The data prove that turning on BRAM ECC (several are used, mostly for buffering) is helpful, see the blue triangles. Most helpful (the brown diamonds) was full TMR with separate clocks although that requires a lot more resources and timing care. However, the data set also demonstrates that “safe” state machine technology harms robustness in reconfigurable FPGAs, as expected; see the black star point. Further, TMR with a single clock also hurts robustness and consumes resources (pink upside-down triangles).




Figure 6. Mitigations, like removing the PLL or adding TMR, deliver less upset-ability

To put the mitigation investigations into context, the unmitigated RapidIO would expect a bit error about twice a year in GEO which is already a very low rate. Adding full TMR is expensive in resources, but offers an order of magnitude improvement, if needed. Incorporating the enhanced BRAM ECC feature into the IP, gives about a factor of five improvement while consuming very little in the way of extra resources.

5.5 Comparing Raw and Protocol Test Results

A presentation is available [17] with a detailed comparison of the Aurora results and the raw MGT results. It shows that, as expected, adding a protocol nets robustness on top of the inherent low sensitivity of the MGTs themselves. In addition, it is possible to enhance a protocol's availability with the addition of a small, but smart supervisor to recognize problems and recover more quickly and transparently as BYU has recently demonstrated for the Aurora protocol.

A comparison of the RapidIO protocol's test results and one of the raw MGT results is shown in Figure 7 below.




Figure 7. Comparing RapidIO errors seen to the reset result of the raw MGT data

6 INPUT/OUTPUT (I/O)-RELATED BLOCKS

The bidirectional Input/Output Blocks (IOB) have optional Single Data Rate (SDR) or Double Data Rate (DDR) registers, serializers and deserializers (SERDES) enabling support for many industry input/output (I/O) standards, selectable output drive strengths and digitally controlled output impedance (DCI), see UG190 [4]. The IOB registers have been implemented with the same RHBD dual node latch as the CLB registers but without the addition of SET filters. Triple-modular redundancy (TMR) was implemented in the DCI control circuitry to improve that feature's upset characteristics to near immeasurability.

6.1 LVCMOS I/Os

The test methodology, designed and implemented by George Madias of Boeing, is based on an asymmetrical loopback where one DUT input ties to three outputs. Thus, it is capable of sensing and counting individual bit errors and small bursts of errors and identifying the cause as either an input or output IOB. In addition it separately measures more troublesome errors that affect many bits and persistent control errors. Routing upsets in the intentionally short loopback paths are very rare due to the upset-hard configuration cells, but would show up as longer duration errors that are fixed by configuration scrubbing. The LVCMOS voltage setting used in this testing was 3.3V although the results are not expected to be very sensitive to the I/O drive voltage. Testing was done with 28 inputs each connected to four outputs, resulting in a total of 112 outputs distributed over four banks of IOBs. Testing was run at 3 and 25 MHz with only a small difference in the results, but, surprisingly especially for the registered case, the 3MHz data appears higher than the 25 MHz data fairly consistently.

The combined results for the checkerboard pattern are shown in Figure 8 and the Weibull fit parameters are tabulated in Table 7 below. In addition to separating input hits from output hits on individual pins, there are less frequent hits that affect an entire bank and “global” hits that affect all four banks. These are all brief outages that self-recover. Note that input hits are the least likely while the other three categories are about equal for heavy ions.

Table 7. Weibull Fit Parameters for Virtex-5QV LVCMOS IOB Upsets from Heavy Ions

	Weibull Parameters			
	Limit (cm ² /IO)	Onset (MeV-cm ² /mg)	Width -	Power -
Unregistered, Output	3.05x10 ⁻⁸	0.5	7	1.00
Unregistered, Input	8.10x10 ⁻⁸	0.5	5	1.00
Registered, Output	2.16x10 ⁻⁵	1.0	9998	1.25
Registered, Input	1.22x10 ⁻⁵	0.70	9998	1.15
	(cm ² /bank)			
Unreg. Bank-wide Transients	1.06x10 ⁻⁷	0.40	2	1.00
Reg. Bank-wide Transients	6.29x10 ⁻⁵	0.50	9998	1.00
	(cm ² /device)			
Unregistered, Globals	3.30x10 ⁻⁸	0.7	3	1.00
Registered, Globals	1.68x10 ⁻⁷	0.621	30.25	0.884




Figure 8a. Four Upset Types and their Cross Section Curves for Unregistered IOBs [18]




Figure 8b. Four Upset Types and their Cross Section Curves for Registered IOBs [18]

The legends of Fig. 8 include GEO rates in units of upsets per bit-day for the input and output cases, in upsets per bank-day for bank transients, and in upsets per device-day for global I/O transients. Note that Boeing used CREME96 [8] parameters that deviate from the XRTC norm in a two ways: (1) x and y dimensions are set as the square root of the cross section at an LET of 80 and (2) the sensitive depth is set to 2 microns with a half micron funnel depth.

6.2 LVDS I/Os

The same test methodology used in the single-ended LVCMOS testing was adapted for 2.5V LVDS differential pairs where each input or output now requires two pins in order to add common-mode noise immunity and increase I/O performance and margin. Apparatus pin limitations restricted the LVDS DUT design to 12 inputs (now differential pairs of pins) and 44 outputs in four banks.

Heavy ion tests covering a full range of LETs were accomplished in September 2012 at the Texas A&M Cyclotron Institute following a preliminary run at Berkeley's BASE Facility in August. Two designs were irradiated: unregistered and registered. In both cases and as in the LVCMOS case above, an input is connected to multiple outputs in order to allow discrimination between hits on an *input* IOB which show up with the same signature on a specific set of multiple outputs simultaneously versus hits on an *output* showing up only on that output. The registered design takes advantage of the built-in IOB flip-flops and thus includes a clocking network in the DUT whereas the unregistered design has no clocking or registration in the DUT. Most of the testing was done with an alternating pattern of data (checkerboard) flowing at 100 Mb/s through the IOBs although the unregistered design was clocked at 25 MHz for some runs to look for frequency effects.

The observed single-event effects are all transitory affecting one or more clock cycles of data. As in the LVCMOS case, four main signatures were evident: SETs on (1) individual OUTPUTs, (2) individual INPUTs, (3) all I/O's on a BANK, and (4) all I/Os tested, dubbed a GLOBAL hit. In addition, two other expected signatures were seen: (1) SEFIs and (2) configuration hits; these are fixed by ConfigMon via corrective SEFI detection and scrubbing, respectively. Note that 'corrective' means issuing the startup command for SU SEFIs and reconfiguring for the others. Figure 9 contrasts the normalized (per I/O) cross section v. LET for the unregistered inputs and outputs at 100 MHz; it is clear that, in the unregistered case, the inputs are less susceptible by about an order of magnitude, similar to the LVCMOS result of the previous section.




Figure 9. Heavy Ion Results for Unregistered Virtex-5QV LVDS IOB Transients.

The SET durations follow a mainly bimodal distribution where one mode, dubbed NORMAL, has about a 2/3 microsecond mean width with a high tail out to about 2 μ s and the other mode is one clock cycle long, dubbed SINGLES, independent of frequency. The input and output data points in Figure 9 for the unregistered case and in Figure 10 for the registered case include both modes. The underlying fractions are observed to vary as a function of LET; as recorded in Table 8; the proportion of singles grows with increasing LET for outputs and are practically all that's observed for inputs.

Table 8. Fraction of Single Bit Error Transients in Virtex-5QV LVDS IOBs

LET	Outputs		Inputs	
	Unreg	Reg	Unreg	Reg
0.89	0.0%	0.0%	--	--
4.00	0.0%	2.3%	--	--
9.00	0.0%	18.0%	--	100.0%
23.3	0.0%	45.8%	100.0%	66.7%
60.7	4.6%	77.6%	100.0%	90.5%




Figure 10. Heavy Ion Results for Registered Virtex-5QV LVDS IOB Transients.

Weibull fits are also shown in the Figures and tabulated in Table 9. It should be noted that fitting Weibull curves to data that doesn't "saturate," that is, that rises sharply at higher LETs as this data does, is usually difficult and unsatisfactory. The fits given here address that problem in a conservative way while weighting heavily toward fitting better in the LETs that contribute most heavily to the space rate. For example, in the unregistered output fit, the curve misses the "knee" points on the high side and this LET range contributes less than 20% to the overall rate anyway.

Table 9. Weibull Fit Parameters for Virtex-5QV LVDS IOB Upsets from Heavy Ions

	Weibull Parameters			
	Limit (cm ² /DSP)	Onset (MeV-cm ² /mg)	Width -	Power -
Unregistered, Outputs	7.0x10 ⁻⁷	0.82	150	1.2
Unregistered, Inputs	6.0x10 ⁻⁸	0.82	150	1.2
Registered, Output Singles	8.0x10 ⁻⁶	0.82	240	2.4
Registered, Input Singles	6.5x10 ⁻⁶	0.82	230	2.1
(cm²/bank)				
Bank-wide Transients	1.3x10 ⁻⁶	0.82	240	2.4
(cm²/device)				
Unregistered, Globals	2.0x10 ⁻⁷	0.82	150	1.2
Registered, Globals	6.85x10 ⁻²	0.2	370	4.3

A different fitting problem is encountered in the unregistered input data of Figure 9- sparse data- so the LET threshold is not well defined. No input transients were encountered at the LETs below 23 MeV-cm²/mg and, thus, the logarithmic plot shows only the tops of the 95%-confidence error bars. The conservative approach taken here is assuming the same LET threshold, indeed even the same shape, as better defined unregistered output transients fit.

Upon encountering the ‘sharply rising’ fitting problem, sometimes a more satisfactory approach is combining two Weibull fits representing two different underlying responses. This approach was taken to the registered data of Figure 10. It makes sense that the registered susceptibilities would include the unregistered response plus whatever the registers and their clock and control trees add. Noting that SINGLES contribute very little to the unregistered results and that, for both inputs and outputs, the NORMAL registered results track well with the total unregistered results adding those two fits to get the output curve shown in Figure 10 which fits the data well.

Bank transient data and fit are shown in Figure 11 for both the registered and unregistered case. These are very rarely seen in the testing and the small amount of data suggests that there is little or no difference between the two.




Figure 11. Heavy Ion Results for Virtex-5QV LVDS IOB Full Bank Transients.

The data and fits for the global transient, where all (or almost all) I/Os are affected, are shown in Figure 12. While unregistered IOBs exhibit very few global I/O SETs, the registered test was quite a bit more susceptible, especially at high LET. However, this

should be quite rare in space; for example, in GEO, the rate is one every few centuries. Note that the test methodology cannot distinguish whether globals affect only outputs or only inputs or both.




Figure 12. Heavy Ion Results for Virtex-5QV LVDS IOB Global (All Banks) Transients.

For proton testing, the test was ported to the Virtex-5QV version of the motherboard and, because the MGT DUT board was used, its Service FPGA also was replaced by a Virtex-5QV FPGA; this reduced test apparatus upsets to insignificance. Proton testing of the LVDS I/Os was conducted during the November and December 2012 test trips. Testing was done at three frequencies with two beam energies and two DUT designs: with and without the IOB flip-flops (“registered” and “unregistered” cases). Results for all the types of hits are shown for the two cases in Figures 13 and 14, respectively with the fit parameters given in the first two rows of Table 10.

The vast majority of these are hits in outputs. In the unregistered case, out of 119 hits observed, all were the output type and had a duration between 70 and 860 ns except one that was 1.32 us. In the registered case, the other three types were observed but too rarely to obtain reasonable statistical significance; out of 271 total events, eleven were the input type, one was bank-wide, and five were global. In addition, about a third of the I/O hits, including all the input hits, had a new signature duration of exactly one clock cycle. This was likely caused by a clock transient on the I/O flip-flop being used to register the data; a direct upset of the flip-flop is unlikely as these are hardened like the CLB flip-flops but without the optional SET filter. Subtracting the unregistered test results from the registered yields the I/O flip-flop susceptibility: there is no difference at 18.8 MeV (with

admittedly low statistics), but a little more than half the 64 MeV proton effects come from the flip-flops. Thus, it is sensible to model this data set as the sum of two separate phenomena with two Weibull fits, one for the flip-flops with a higher LET threshold and one for the rest of the IOB as is shown in Figure 15. The rest of the IOB (minus the flop-flop) should be the same as the unregistered case of Figure 13 and the data is not inconsistent with that hypothesis. Table 10 includes these additive fits in the last two rows as an alternative to the combined fit for registered IOBs in the preceding row.

Table 10. Weibull Fit Parameters for Virtex-5QV LVDS 2.5V IOB Upsets from Protons

	Weibull Parameters			
	Limit (cm ² /IOB)	Onset (MeV)	Width -	Power -
Unregistered, All Events	3.5x10 ⁻¹⁴	5.	12	1
Registered, All Events	1.5x10 ⁻¹³	5.	80	1
Registered IOB only	3.5x10 ⁻¹⁴	5.	12	1
Registered FF only	4.5x10 ⁻¹⁴	30.	6	1




Figure 13. Proton Results and Weibull Curve for All Effects on Unregistered IOBs.




Figure 14. Proton Results and Weibull Curve for All Effects on Registered IOBs.




Figure 15. A Proposed Better Fit for the Proton Results for All Effects on Registered IOBs

6.3 IOSERDES

The IOSERDES feature takes advantage of the higher speed 2.5V LVDS I/O pins by adding a 4-bit serializer to outputs and a 4-bit deserializer to inputs. The test methodology for these IP blocks is similar to the I/O testing described in the previous two sections with a few small differences. Each input's deserializer loops back to a single output's serializer in the DUT. This means that the signature of input and output hits are the same and can't be separated. This allowed raising the number of I/O pairs under test to 64, that is, 32 input differential pairs and 32 output differential pairs are under test. Another enhancement is separating the bits-in-error counters into separate counters for zeros-in-error and ones-in-error. In combination with the alternating or checkerboard pattern, this enhancement allows discrimination between a hit that causes the data to be high or low (only one counter counts) versus a synchronization or clock hit (where both counters increment while the condition persists). The heavy ion results indicate that these two mechanisms have different cross section curves, but end up making about the same contribution to the GEO rate: $\sim 10^{-6}$ per bit day or about once every 2500 years.

The DUT design has 32 channels using four I/O banks with each channel having the following characteristics: (1) input and output pins in the same I/O bank, (2) input pin connected to a ISERDES serial to 4-bit parallelizer, (3) with a short routing on 4-bit bus to (4) a OSERDES serializer connected to an output pin. Operation of the IOSERDES heavy-ion tests used a checkerboard pattern at 3 MHz and relied on pre-production software not yet capable of automatically removing half-latches nor were they removed manually. Thus, the non-persistent error modes seen were similar to a variation of the LVCMOS IOB registered testing (not reported here) that intentionally did not remove half-latches [18]. Three other error signatures were measured: (1) so-called "1 or 0" events where the output is stuck at a particular value, (2) "1 and 0" events where both values are erroneous corresponding to clock hits that got the input and output streams out of step with each other, and (3) I/O bank upsets where all channels on a given bank were not operating properly. Most of the "1 or 0" stuck events and the "1 and 0" out-of-alignment events were fixed by resetting the ISERDES and OSERDES blocks. Note that no global-type I/O outages were observed during this test although the IOB tests certainly indicate that they are possible.

Table 11a. Weibull Fit Parameters for Virtex-5QV IOSERDES Upsets from Heavy Ions

	Weibull Parameters			
	Limit (cm ² /IO)	Onset (MeV-cm ² /mg)	Width -	Power -
'1 or 0' - stuck	6.0x10 ⁻²	0.77	10,000	1.997
'1 and 0' - clock hits	6.46x10 ⁻⁸	0.49	13.9	0.997
	(cm ² /bank)			
Bank Upsets	5.82x10 ⁻⁵	0.19	10,000	0.918




Figure 16a. Heavy Ion Results and Weibull Curve for the IOSERDES feature.




Figure 16b. Proton Results and Weibull Curve for the IOSERDES feature.

The proton test on the IOSERDES feature was conducted as part of the December 2012 test trip to the UC-Davis cyclotron. Data was taken at two proton energies and operating at two frequencies. The results are plotted in Figure 16. The sparse data set is consistent with a two component response like that shown for registered LVDS in Figure 15. Again, the lower energy threshold element would be the basic unregistered IOB response of Figure 13 combined with, in this case, the response of the serializer/deserializer.

Table 11b. Weibull Fit Parameters for Virtex-5QV IOSERDES Upsets from Protons

	Weibull Parameters			
	Limit (cm ² /IO)	Onset (MeV-cm ² /mg)	Width -	Power -
All Events	9.0x10 ⁻¹⁴	5.	52	1
IOB only	3.5x10 ⁻¹⁴	5.	12	1
SERDES only	3.5x10 ⁻¹⁴	30.	6	1

The postulated two Weibull response fit is included as the last two rows of Table 11b; added together these give the total response measured. Alternatively, the data can be fitted more conventionally with a single Weibull curve and those parameters are given in the first row of Table 11b and shown graphically in Figure 16b.

6.4 IODELAY

The IODelay feature is individual I/O line de-skew capability implemented as a programmable delay line. This feature is a major convenience for mating to high speed bus interfaces, usually memories like QDR or DDRx, as very tiny timing changes are programmable and can compensate for board trace length differences and/or memory device timing variations.

The experiment measures SEE on both the per-pin IODelay feature and the regional IODelay Controllers (there are a total of 28 of these in the XQR5VFX130 as well as its commercial counterpart, the XC5VFX130T). The experiment design by George Madias of Boeing uses 58 output pins with IODelay instantiated in a ring oscillator along with 13 IODelay Controllers. In the first XRTC heavy ion tests, both phenomena show a kink in the cross section vs. LET curves as can be seen in Figure 17a and 17b; this data was taken in July 2009 at the Texas A&M cyclotron. Such a kink usually indicates two underlying susceptible elements, one with a low threshold LET and a lower cross section at high LET. It's a bit unusual for this to be so clear in the data and the underlying elements causing this, if indeed the phenomenon is real and not an experimental artifact, have not been identified. One reason to suspect an experimental artifact, like a malfunctioning DUT or a dosimetry issue, is that the data points at LET = 25 and 35 were taken with a different DUT (s/n: 33) and ion (25 MeV/amu Kr) than the others (s/n: A4114 and A4415 and 25 MeV/amu Ar, Xe, and Ne and 15 MeV/amu Au). Later testing of QDR SRAM interface IP that incorporates the IODelay tends to confirm the "experimental artifact" explanation of the kinks.




Figure 17a. Alternative Fits of the Heavy Ion Data for IODELAY Upsets.




Figure 17b. Alternative Fits of the Heavy Ion Cross Section for Persistent Problems Caused by Upsets to the IODELAYCONTROL Circuitry.

Table 12a. Weibull Fit Parameters for Virtex-5QV IODELAY Upsets from Heavy Ions

	Weibull Parameters			
	Limit (cm ² /IO)	Onset (MeV-cm ² /mg)	Width -	Power -
IODELAY	6.0x10 ⁻⁴	0.1	400	1.5
IODELAYCONTROL	2.5x10 ⁻⁴	0.1	300	1.9

Table 12b. Weibull Fit Parameters for Virtex-5QV IODELAY Upsets from Heavy Ions

	Weibull Parameters			
	Limit (cm ² /IO)	Onset (MeV-cm ² /mg)	Width -	Power -
IODELAY	4.0x10 ⁻²	1.89	9999	1.38
IOController (per 13 controllers)	3.1x10 ⁻¹	0.94	9999	2.05

Table 13. GEO Rates for Virtex-5QV IODELAY Upsets
Comparing the Results from the Weibull and Double Exponential Fits

	Weibull		Exponential	
	Best Est.	Worst Case	Best Est.	Worst Case
IO Delay	7.7e-5 / IODelay-Day	1.1e-4 / IODelay-Day	2.8e-5 / IODelay-Day	4.0e-5 / IODelay-Day
	2.8 / IODelay-Century	~4 / IODelay-Century	~1 / IODelay-Century	~1.5 / IODelay-Century
IO Controller	9.1e-6 / IOController-day	1.8e-5 / IOController-day	7.14e-6 / IOController-day	1.34e-5 / IOController-day
	~3 failures per IOController per 1000 years	~7.5 failures per IOController per 1000 years	~2.6 failures per IOController per 1000 years	~5 failures per IOController per 1000 years

Note that a proton experiment on the IODelay features was not undertaken.

As seen in Table 13, the Weibull fit consistently yields about a factor of three higher rates than the exponential fit for the IODelay hits and may therefore be overestimating the GEO rate significantly. For the persistent hits attributed to the IODelay Controllers the difference between the two fits is less dramatic, in the range of 15-35%, so the Weibull fit is not so severely conservative.

7 CLOCK MANAGEMENT (CMT)

While it is possible to directly use clocks brought from off chip, there are some very powerful and flexible clock management resources built into the silicon on Virtex-5 family devices, including the Virtex-5QV. They are provided as Clock Management Tiles (CMTs) with each tile having 3 separate clocking resources consisting of two Digital Clock Managers (DCMs) and one Phase Locked Loop (PLL). The Virtex-5QV includes 6 CMTs and therefore 12 DCMs and 6 PLLs.

The Digital Clock Manager (DCM) blocks provide clock dividers and multipliers, clock frequency synthesis, phase shift and de-skew capabilities, and dynamic reconfiguration. The Phase Lock Loop (PLL) blocks provide clock distribution delay compensation, clock multiplication/division, coarse-/fine-grained clock phase shifting, and input clock jitter filtering. Details on these CMT components and how to use them are given in Chapters 2 and 3 of UG190 [4]

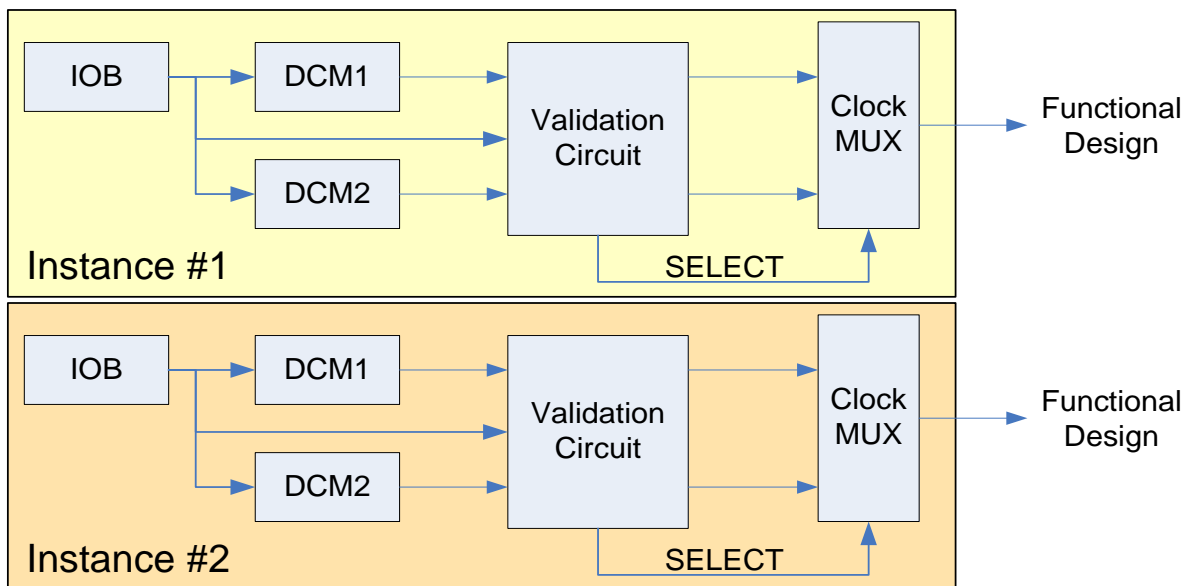


Figure 18. DUT Design Block Diagram [after Ref. 18 & 19]

A common test method was used for both DCMs and PLLs that was able to detect the main upset and transient modes to which they are susceptible. A block diagram of the DUT design by George Madias of Boeing is shown in Figure 18 for DCMs; the PLL version is the same with PLLs replacing DCMs. Thus, the Functional Monitor for both tests is identical. Two instances of the same basic structure not only double the target area which makes beam use more efficient, but also allow the observation of events that affect both instances. Each instance consists of a primary and secondary clocking source. If the primary is upset, the validation circuit will switch them while resetting the “switched out” source. If the secondary was affected by the same event, then the validation circuit will switch again and reset it as well. For a really long transient, switching and resetting the clock sources may occur multiple times. The inputs for the two instances are in two separate I/O banks and are therefore not susceptible to an IOB bank transient, only a global IOB transient. Consequently, there are four output

categories that should be counted separately, single or multiple switches occurring in one or both instances.

There are four error signatures extracted from the FuncMon logs that are presented below; they correspond to the above four categories as follows:

- 1) Single switches on only one instance are dubbed “unique singles;”
- 2) All “single switches” are counted regardless of what’s happening on the other instance, so the unlikely case of both instances experiencing single switches is derived by subtracting “unique singles” from “single switches;”
- 3) Multiple switches on both instances are called “both switch multiple times;”
- 4) Somewhat confusingly, multiple switches on only one instance are labeled “multiple switches” instead of “unique multiples.”

Note that signature #4 does not include signature #3, but signature #2 does include #1.

To some extent, signatures and their causes can be correlated. Multiple switches are the result of hitting something the primary and secondary of an instance have in common while hitting both instances requires a hit to a cross-instance shared resource. It is clear, for example, that IOB globals will cause “both switch multiple times” and that for the duration of an IOB bank transient, “multiple switches” on one instance will occur. Also, any hit on the common routing from the pin to the clock source inputs or from the pin to the validation circuit will cause “multiple switches” until scrubbed; see Figure 19. The bottom line is that “single switches” really represent the susceptibility of the individual clock source (DCM or PLL) itself.




Figure 19. Single Instance Multiple Switches Hits Affect Both PLL and DCM about Equally [Ref 19, Fig7 and Ref 20 Slide#10]

Additionally, the Functional Monitor observes a running counter for each instance's output. When one or both counters stop for about five seconds, FuncMon and the validation circuit are manually reset and a "reset" event is counted. This gives enough time for events caused by a configuration upset to be scrubbed out as ConfigMon accomplishes a readback and scrub cycle in less than one second. It is important to note that no SEFI-like events were seen in this testing; that is there were no clock outages that could not be reset; thus, reconfiguration was never needed to recover an upset clock source.

Somewhat surprisingly, the results for the two very different types of clock management resources are very similar and the resulting GEO rates are very close. It makes sense that the input results might be similar as the design rules and libraries are common and, of course, the common test methodology assures that the outside sources of SETs (like an IOB global) are the same. So the similarity in input response shown in Fig. 19 is not too surprising. To the extent that "multiples" and "both multiples" are measures of the susceptibility of the validation circuitry, it's also not surprising that they are the same, regardless of the type of clock source. Three of the four signature curves could have had their data combined to be fitted with the same Weibull parameters and fitting separately only yields tiny differences. Only in "unique singles" (the black asterisks) do the two types of clocking sources show a clear difference with the PLL significantly lower.

7.1 Digital Clock Management (DCM) Blocks




Figure 20a. Susceptibility to Four DCM Error Signatures [courtesy of E. Miller, Boeing]
Note in the legend “Bank” means “instance” and the data is 2x low.

The four types of output upsets caused by heavy ion bombardment have about the same low LET threshold, but the cross section at high LET shows about an order of magnitude spread in susceptibility. For the fits shown, small variations in the threshold region almost exactly offset the high LET differences resulting in rates for all four types that are about the same (within a factor of 3) for GEO. It is sensible that “unique singles” make up all or almost all of the “single switches” as it is hard to envision how both instances’ primaries can be hit without affecting one or both of the secondaries. It is important to note that this figure was generated with an erroneous calibration factor and the data and fits should be 2x higher.

The legend of Fig. 20a includes GEO rates in units of upsets per DCM-day. Note that Boeing used CREME96 [8] parameters that deviate from the XRTEC norm in a two ways: (1) x and y dimensions are set as the square root of the cross section at an LET of 80 and (2) the sensitive depth is set to 2 microns with a half micron funnel depth. Comparing the rates given is valid, but the actual rates given are inaccurate because they include the erroneous calibration factor.

As might be expected from the low heavy ion susceptibility, the proton SEU sensitivity is also quite low. Extensive use of the clocking resource tests for the SEFI campaign has

resulted in a solid foundation with good statistics. As a consequence, the 64 MeV measurements have very tiny 90% error bars and the measurement at 18.8 MeV is almost as good. Interestingly, under proton irradiation some of the signatures seen with heavy ions become disproportionately rare. For instance, a large majority of the switches logged under proton irradiation are the simple single type with multiples –which are almost always just double switches- too infrequent to break out separately (more than order of magnitude fewer). Also shown is a second category: the longer outages of a tenth of a second or more, which either a) self-recover eventually or b) recover after a scrub cycle or c) require resetting the validation circuitry; these are all lumped in with “resets” as they are individually too rare and also because it is difficult to definitively discriminate between their signatures in the event logs. Most of them are probably a result of problems with the validation circuit.




Figure 20b. Proton Susceptibility Lumped into Two Categories of DCM Error Signatures.

Weibull parameters for the heavy ion fits of the DCM data are given in Table 13a; similarly the proton fitting parameters used for the curves in Figure 20b are listed in Table 13b.

Table 13a. Weibull Fit Parameters for Virtex-5QV DCM Upsets from Heavy Ions

	Weibull Parameters			
	Limit (cm²/DCM)	Onset (MeV-cm²/mg)	Width	Power -
DCM - single	3.86x10 ⁻³	0.25	9,998	1.25
DCM - unique single	7.82x10 ⁻⁴	0.50	9,998	1.00
DCM - multiple	1.73x10 ⁻²	1.85	9,998	1.36
DCM - both multiple	9.72x10 ⁻²	1.23	9,998	1.49
Loss of Lock	1.22x10 ⁻⁴	1.97	9,998	0.85

Table 13b. Weibull Fit Parameters for Virtex-5QV DCM Upsets from Protons

	Weibull Parameters			
	Limit (cm²/DCM)	Onset (MeV)	Width (MeV)	Power -
DCM -switch	3.5x10 ⁻¹²	5.0	13	1.0
Validation circuit - reset	4.0x10 ⁻¹³	5.0	9	1.0

7.2 Phase-Locked Loop (PLL) Blocks

The same four types of responses as with DCMs are seen for PLLs under heavy ion bombardment. Although for PLLs, the fits yield about an order of magnitude spread in rates in GEO, they sum to about the same combined rate. For both DCMs and PLLs the “both switch multiple times” category is the predominant response to heavy ions, as can be clearly seen in Figure 20a for DCMs and 21a for PLLs. Note that this figure has an erroneous calibration factor and the data and fits should be 2x higher.




Figure 21a. Susceptibility to Four PLL Error Signatures [courtesy of E. Miller, Boeing]
Note in the legend “Bank” means “instance” and the data is 2x low.

The legend of Fig. 21a includes GEO rates in units of upsets per PLL-day. As a reminder, Boeing’s choice of CREME96 [8] parameters deviates from the XRTC norm in a two ways: (1) x and y dimensions are set as the square root of the cross section at an LET of 80 and (2) the sensitive depth is set to 2 microns with a half micron funnel depth.

The proton responses of the PLL circuitry is also very similar to that of the DCM presented above and again that response is dominated by the least intrusive category, that of single switches on only one instance. Figure 21b compares the fit and the data for both switches and resets, the same categorization as used for the DCM proton results of Figure 20b.




Figure 21b. Susceptibility to Two PLL Error Signatures Observed in Proton Tests

The susceptibilities of the Virtex-5QV PLLs to heavy ions and to protons are summarized as fitted Weibull parameters in Tables 14a and 14b, respectively.

Table 14a. Weibull Fit Parameters for Virtex-5QV PLL Upsets from Heavy Ions

	Weibull Parameters			
worse case	Limit (cm ² /PLL)	Onset (MeV-cm ² /mg)	Width -	Power -
PLL - single	1.21x10 ⁻⁵	0.50	56.9	1.68
PLL - unique single	2.20x10 ⁻⁶	0.35	29.7	1.70
PLL - multiple	1.89x10 ⁻⁵	0.33	32.9	2.17
PLL - both multiple	6.44x10 ⁻⁵	0.70	43.5	1.85
Loss of Lock	1.17x10 ⁻⁷	1.45	33.7	1.33

Table 14b. Weibull Fit Parameters for Virtex-5QV PLL Upsets from Protons

	Weibull Parameters			
worse case	Limit (cm ² /PLL)	Onset (MeV)	Width -	Power -
PLL -switch	2.5x10 ⁻¹²	7.0	5	1.0
Validation circuit - resets	6.0x10 ⁻¹³	7.0	20	1.0

There are two final comments on the DCM and PLL results. First, the testing did observe the state of each clock primitive's 'Lock' signal. As was seen in previous families, the beam does cause it to change state, but in a way that is not connected with the actual health of the clock or its lock status. This signal is really only meaningful when it first transitions from 'not locked' to locked after either configuration or de-asserting reset. Although included in the Weibull parameter tables, knowing the susceptibility of the lock signal isn't really useful for anything.

Second, the test methodology suggests a dual clocking mitigation scheme. Under such a scheme, the single switches would be almost transparent to system operation and multiple switches would recover a usable clock as quickly as possible so that only the less frequent longer self-recovering outages and those requiring reset would be intrusive; based on this data, an increase in robustness of an order-of-magnitude or more can be expected in proton-rich environments. Further circuitry to issue a reset as needed would be a natural adjunct; reset circuitry would significantly increase availability by reducing the durations of single-event clock outages. In other words, the "reset" cross section noted is really a measure of the failure susceptibility of this mitigation scheme.

8 SPACE UPSET RATES FOR SELECTED ORBITS

Dynamic SEU characterization of thirteen types of silicon blocks or “architectural features” or sub-features were summarized in the preceding Sections and Weibull parameters for thirty-four upset phenomena were tabulated. These allow the calculation of expected upset rates for any orbit. Some representative orbits are tabulated here in this final Section as concrete examples. For a few of the upset phenomena, no proton data was taken and only the heavy ion contribution to the rate is listed. There are several methods available for predicting proton rates from heavy ion measurements, but none has broad acceptance and all have known problems and the authors have opted not to choose here. Clearly, additional proton testing would be the best answer for those phenomena.

Most of the architectural features are not upset hardened themselves although the hardening of the configuration cells which select and program their features is a great boon. The unhardened features tend to have fairly low LET thresholds counterbalanced by low cross sections at high LET, resulting in low or moderate space environment upset rates. Proton reaction products will cause upsets in these; thus, GEO is, as usual, a best-case environment. As the orbital altitude is lowered into MEO or LEO, the additional upsets from trapped protons more than offset the reduction in GCR heavy ion upsets.

Three of the architectural features incorporate upset hardening or mitigation into the hardware- BRAM ECC, dual-node User Flip-Flops along with their associated SET filters, and the DCI capability of the IOBs. For these, their individual upset stories are a bit more complicated. Normally, the rate at which BRAM errors are seen is strongly dependent on the peak error rates. However, in this case, there are some single points of failure in the ECC and the rate of upset of these fail points dominates. Likely they are fairly proton susceptible so that GEO is again a best-case environment.

The dual node hardening of the User Flip-Flop makes them nearly immune to upsets induced by proton reaction products. Further, the SET filters, if used, are especially effective at benignly soaking up the shorter pulses associated with low LETs and proton reaction products. Thus, GEO is pretty close to the worst-case environment for upset-hardened elements like the User Flip-Flops; that is, for most orbits, their rates actually lower than in GEO and the shielding effect of the earth’s magnetic field increases lowers heavy ion rates more the additional proton component.

Table 15a. CREME96 Calculated Orbital Upset Rates

User Flip-Flops Cells, 200 MHz, Virtex-5QV						
Upsets/Device-Day with Proton Portion in parentheses, assuming all 81920 are used						
Quiet Solar Minimum Conditions, AP8MIN, 100 mils Al						
Orbit	LEO-low	LEO-mid	LEO-high	POLAR	GPS	GEO
Altitude (km)	500	800	1200	833	20,200	36,000
Inclination	51.6°	22.0°	65.0°	98.7°	55.0°	0.0°
Filters ON	1.34x10 ⁻⁴ (37%)	4.60x10 ⁻⁴ (95%)	1.33x10 ⁻³ (84%)	5.15x10 ⁻⁴ (50%)	8.19x10 ⁻⁴ (0.97%)	8.93x10 ⁻⁴ (0.95%)
Filters OFF	2.28x10 ⁻³ (32%)	6.87x10 ⁻³ (95%)	2.06x10 ⁻² (81%)	8.74x10 ⁻³ (43%)	1.60x10 ⁻² (0.74%)	1.72x10 ⁻² (0.73%)

Table 15b. CREME96 Calculated Orbital Upset Rates**DSPs, 25 MHz, Virtex-5QV**

Upsets/Device-Day with Proton Portion in parentheses, assuming all 320 are used
 Quiet Solar Minimum Conditions, AP8MIN, 100 mils Al

Orbit	LEO-low	LEO-mid	LEO-high	POLAR	GPS	GEO
Altitude (km)	500	800	1200	833	20,200	36,000
Inclination	51.6°	22.0°	65.0°	98.7°	55.0°	0.0°
Multiply	1.60x10 ⁻² (48%)	6.84x10 ⁻² (96%)	1.92x10 ⁻¹ (91%)	5.97x10 ⁻² (66%)	5.85x10 ⁻² (1.8%)	6.21x10 ⁻² (1.8%)
Add	1.39x10 ⁻² (48%)	6.11x10 ⁻² (96%)	1.67x10 ⁻¹ (91%)	5.20x10 ⁻² (66%)	5.05x10 ⁻² (2.2%)	5.36x10 ⁻² (2.2%)
Accumulate	5.40x10 ⁻³ (51%)	2.52x10 ⁻² (97%)	6.82x10 ⁻² (92%)	2.08x10 ⁻² (68%)	1.92x10 ⁻² (2.4%)	2.03x10 ⁻² (2.4%)

Table 15c. CREME96 Calculated Orbital Upset Rates**MGTs (GTXs), 3.25 G/s, Virtex-5QV**

Heavy Ion Upsets/Device-Day*, assuming all 18 are used
 Quiet Solar Minimum Conditions, AP8MIN, 100 mils Al

Orbit	LEO-low	LEO-mid	LEO-high	POLAR	GPS	GEO
Altitude (km)	500	800	1200	833	20,200	36,000
Inclination	51.6°	22.0°	65.0°	98.7°	55.0°	0.0°
Bit Errors	1.35x10 ⁻³	4.28x10 ⁻⁴	2.97x10 ⁻³	3.39x10 ⁻³	9.65x10 ⁻³	1.03x10 ⁻²
Transmitter Loss-of-Link	1.35x10 ⁻⁴	4.17x10 ⁻⁵	3.03x10 ⁻⁴	3.49x10 ⁻⁴	1.01x10 ⁻³	1.08x10 ⁻³
Receiver Loss-of-Link	2.84x10 ⁻⁴	8.90x10 ⁻⁵	6.30x10 ⁻⁴	7.22x10 ⁻⁴	2.07x10 ⁻³	2.21x10 ⁻³

*-Portion of the total rate due to heavy ions only. Proton susceptibility not measured, so proton rate not included.

Table 15d. CREME96 Calculated Orbital Upset Rates**IOBs, LVDS, 100 MHz, Virtex-5QV**

Upsets/Device-Day, Proton Portion in parentheses, assuming all 418 pairs are used, half as inputs
 Quiet Solar Minimum Conditions, AP8MIN, 100 mils Al

Orbit	LEO-low	LEO-mid	LEO-high	POLAR	GPS	GEO
Altitude (km)	500	800	1200	833	20,200	36,000
Inclination	51.6°	22.0°	65.0°	98.7°	55.0°	0.0°
Unregistered	5.49x10 ⁻⁵ (70%)	3.33x10 ⁻⁴ (98.5%)	9.002x10 ⁻⁴ (96%)	2.40x10 ⁻⁴ (81%)	1.40x10 ⁻⁴ (3.8%)	1.50x10 ⁻⁴ (3.8%)
Registered	2.18x10 ⁻⁴ (96.5%)	1.93x10 ⁻³ (99.9%)	4.82x10 ⁻³ (99.6%)	1.13x10 ⁻³ (97.6%)	1.35x10 ⁻⁴ (33%)	1.46x10 ⁻⁴ (32%)
Bank *	2.96x10 ⁻⁸	4.42x10 ⁻⁹	9.06x10 ⁻⁸	1.21x10 ⁻⁷	4.19x10 ⁻⁷	4.62x10 ⁻⁷
Unreg. Global **	1.82x10 ⁻⁸	5.28x10 ⁻⁹	4.28x10 ⁻⁸	5.05x10 ⁻⁸	1.52x10 ⁻⁷	1.63x10 ⁻⁷
Reg. Global ***	1.51x10 ⁻⁵ (32%)	4.49x10 ⁻⁵ (99.0%)	1.47x10 ⁻⁴ (75%)	7.73x10 ⁻⁵ (33%)	1.90x10 ⁻⁴ (0.54%)	2.12x10 ⁻⁴ (0.51%)

*-Heavy ions only and assuming all 23 I/O banks are used. No bank events were seen in unregistered proton tests and only one (<0.4%) in the registered tests.

** -Heavy ions only. No global events seen in the unregistered proton tests.

***-Proton fraction estimated, using the actual data, as 5/217 (1.84%) of all proton-induced events (from row 2).

Table 15e. CREME96 Calculated Orbital Upset Rates**IOBs, LVC MOS, 25 MHz, Virtex-5QV**

Heavy Ion Upsets/Device-Day*, assuming all 816 I/Os are used as 408 inputs and 408 outputs

Quiet Solar Minimum Conditions, AP8MIN, 100 mils Al

Orbit	LEO-low	LEO-mid	LEO-high	POLAR	GPS	GEO
Altitude (km)	500	800	1200	833	20,200	36,000
Inclination	51.6°	22.0°	65.0°	98.7°	55.0°	0.0°
Unregistered	2.87×10^{-4}	9.38×10^{-5}	6.27×10^{-4}	7.12×10^{-4}	2.02×10^{-3}	2.14×10^{-3}
Registered	1.19×10^{-5}	3.59×10^{-6}	2.74×10^{-5}	3.19×10^{-5}	9.43×10^{-5}	1.01×10^{-4}
Bank, reg.	8.30×10^{-6}	2.65×10^{-6}	1.83×10^{-5}	2.09×10^{-5}	5.98×10^{-5}	6.35×10^{-5}
Global ***	2.91×10^{-7}	9.40×10^{-8}	6.40×10^{-7}	7.29×10^{-7}	2.08×10^{-6}	2.21×10^{-6}

*-Portion of the total rate due to heavy ions only. Proton susceptibility not measured, so proton rate not included.

**-Registered, unregistered is about a factor of 5 lower using the fits given in Table 7.

***-Registered, unregistered is about 26% lower.

Table 15f. CREME96 Calculated Orbital Upset Rates**IOSERDES, 3 MHz, Virtex-5QV**

Upsets/Device-Day with Proton Portion in parentheses, assuming all 836 are used with half inputs

Quiet Solar Minimum Conditions, AP8MIN, 100 mils Al

Orbit	LEO-low	LEO-mid	LEO-high	POLAR	GPS	GEO
Altitude (km)	500	800	1200	833	20,200	36,000
Inclination	51.6°	22.0°	65.0°	98.7°	55.0°	0.0°
All*	2.64×10^{-4} (28%)	7.22×10^{-4} (92.9%)	2.15×10^{-3} (79%)	9.42×10^{-4} (41%)	1.73×10^{-3} (0.78%)	1.86×10^{-3} (0.77%)
Bank**	2.35×10^{-5}	7.63×10^{-6}	5.07×10^{-5}	5.72×10^{-5}	1.61×10^{-4}	1.70×10^{-4}

*-'All' means the sum of heavy ion reset-able effects of stucks and clock hits plus proton 'all events.'

Thus, LVDS IOB effects listed above (Table 15d) are included, except for bank or global events, i.e. per pin effects included while per-bank or per-device effects are not.

**- Heavy ions only and assuming all 23 I/O banks are used. No bank events were seen in IOSERDES proton tests.

Table 15g. CREME96 Calculated Orbital Upset Rates**IDELAY, Virtex-5QV**

Heavy Ion Upsets/Device-Day*, assuming all 836 I/Os and 23 IDELAY controllers are used

Quiet Solar Minimum Conditions, AP8MIN, 100 mils Al

Orbit	LEO-low	LEO-mid	LEO-high	POLAR	GPS	GEO
Altitude (km)	500	800	1200	833	20,200	36,000
Inclination	51.6°	22.0°	65.0°	98.7°	55.0°	0.0°
IDELAY	1.01×10^{-2}	3.04×10^{-3}	2.29×10^{-2}	2.67×10^{-2}	7.83×10^{-2}	8.37×10^{-2}
IDELAY Controller	2.74×10^{-5}	7.69×10^{-6}	6.55×10^{-5}	7.81×10^{-5}	2.37×10^{-4}	2.55×10^{-4}

*-Portion of the total rate due to heavy ions only. Proton susceptibility not measured, so proton rate not included.

Table 15h. CREME96 Calculated Orbital Upset Rates**DCMs, Virtex-5QV**

Upsets/Device-Day with Proton Portion in parentheses, assuming all 12 are used
 Quiet Solar Minimum Conditions, AP8MIN, 100 mils Al

Orbit	LEO-low	LEO-mid	LEO-high	POLAR	GPS	GEO
Altitude (km)	500	800	1200	833	20,200	36,000
Inclination	51.6°	22.0°	65.0°	98.7°	55.0°	0.0°
Singles *	1.57x10 ⁻⁴ (70%)	9.52x10 ⁻⁴ (98.4%)	2.57x10 ⁻³ (95.9%)	5.78x10 ⁻⁴ (82%)	3.64x10 ⁻⁴ (4.2%)	3.87x10 ⁻⁴ (4.2%)
Multiples **	2.72x10 ⁻⁵ (40%)	9.78x10 ⁻⁵ (95.7%)	2.86x10 ⁻⁴ (86%)	1.05x10 ⁻⁴ (53%)	1.54x10 ⁻⁴ (0.99%)	1.66x10 ⁻⁴ (0.98%)
Globals ***	8.72x10 ⁻⁶	2.31x10 ⁻⁶	2.13x10 ⁻⁵	2.56x10 ⁻⁵	7.90x10 ⁻⁵	8.50x10 ⁻⁵

*-Using 'unique singles' as it comes out slightly higher than 'singles' which should be equal or greater.

** -Assumes each pair of DCMs (there are six total) hit by same event share a common clock input. Proton fraction estimated from small data set as about 10% of single switches (from row 1 above).

***-Heavy ions only. No global events (both instances, multiples) seen in the proton tests.

Table 15i. CREME96 Calculated Orbital Upset Rates**PLLs, Virtex-5QV**

Upsets/Device-Day with Proton Portion in parentheses, assuming all 6 are used
 Quiet Solar Minimum Conditions, AP8MIN, 100 mils Al

Orbit	LEO-low	LEO-mid	LEO-high	POLAR	GPS	GEO
Altitude (km)	500	800	1200	833	20,200	36,000
Inclination	51.6°	22.0°	65.0°	98.7°	55.0°	0.0°
Singles	5.15x10 ⁻⁵ (79%)	3.49x10 ⁻⁴ (99.1%)	9.45x10 ⁻⁴ (97.3%)	2.37x10 ⁻⁴ (87%)	9.71x10 ⁻⁵ (5.6%)	1.04x10 ⁻⁴ (5.6%)
Multiples *	1.46x10 ⁻⁵ (28%)	3.75x10 ⁻⁵ (92.3%)	1.18x10 ⁻⁴ (78%)	5.15x10 ⁻⁵ (40%)	9.51x10 ⁻⁵ (0.57%)	1.02x10 ⁻⁴ (0.56%)
Globals **	1.04x10 ⁻⁵	2.89x10 ⁻⁶	2.49x10 ⁻⁵	2.97x10 ⁻⁵	9.06x10 ⁻⁵	9.73x10 ⁻⁵

*-Proton fraction estimated as 10% of all single switches (from row 1 above).

** -Heavy ions only. No global events (both instances, multiples) seen in the proton tests.

9 REFERENCES

- [1] G. Swift and G. Allen, *Virtex-5QV Static SEU Characterization Summary*, JPL, Pasadena, CA, July 2012, [online] Available: <http://parts.jpl.nasa.gov/wp-content/uploads/V5QV-Static-SEU-Summary-ReportRevD.pdf> .
- [2] G. Allen, G. Swift, C. Carmichael and the XRTC, *Virtex-4QV Static SEU Characterization Summary*, Sept. 2008, available at <http://parts.jpl.nasa.gov/wp-content/uploads/NEPP07FPGA v4Static.pdf>
- [3] G. Allen, *Virtex-4VQ Dynamic and Mitigated Single Event Upset Characterization Summary*, Jan. 2009, available at <http://parts.jpl.nasa.gov/wp-content/uploads/V4-DM.pdf> .
- [4] Xilinx Documentation: (a) *Virtex-5 FPGA User Guide*, UG190, http://www.xilinx.com/support/documentation/user_guides/ug190.pdf , (b) *Radiation-Hardened, Space-Grade Virtex-5QV FPGA Data Sheet: DC and Switching Characteristics*, DS692, http://www.xilinx.com/support/documentation/data_sheets/DS692.pdf and (c) *Radiation-Hardened, Space-Grade Virtex-5QV Family Overview*, DS192, http://www.xilinx.com/support/documentation/data_sheets/ds192_V5QV_Device_Overview.pdf
- [5] G. Allen et al., “Single-Event Upset (SEU) Results of Embedded Error Detect and Correct Enabled Block Random Access Memory (Block RAM) within the Xilinx XQR5VFX130,” *2010 REDW Record*, available at <http://parts.jpl.nasa.gov/wp-content/uploads/EDAC-final-IEEE-proof.pdf> .
- [6] C. Carmichael and E. Miller, “Characterization & Filtration of Single Event Transient Effects in 65nm CMOS Technology,” SEE Symposium, La Jolla, CA, April 2010.
- [7] G. Swift, SEE Symposium, La Jolla, CA, April 2011
- [8] CREME website: <https://creme.isde.vanderbilt.edu/>
- [9] *Virtex-5 FPGA XtremeDSP Design Considerations User Guide*, UG193, http://www.xilinx.com/support/documentation/user_guides/ug193.pdf
- [10] R. Monreal, “Radiation Test Report, Single Event Effects, Virtex-5QV Field Programmable Gate Array, Digital Signal Processors,” Southwest Research Institute, Austin, TX, Doc. 14520-RATR-03, July 2011. [online] Available: <http://parts.jpl.nasa.gov/wp-content/uploads/14520-RATR-03.pdf>
- [11] R. Monreal and G. Swift, “Upset Manifestations in Embedded Digital Signal Processors (DSPs) due to Single Event Effects (SEE),” *RADECS Data Workshop*, Sept. 2012.
- [12] *Virtex-5 FPGA RocketIO GTX Transceiver User Guide*, UG198, http://www.xilinx.com/support/documentation/user_guides/ug198.pdf .
- [13] R. Monreal on MGTs, available at <http://parts.jpl.nasa.gov/wp-content/uploads/14520-RATR-01Rev-.pdf> .
- [14] R. Monreal, C. Carmichael, and G. Swift, “Single-Event Characterization of Multi-Gigabit Transceivers (MGT) in Space-Grade Virtex-5QV Field Programmable Gate Arrays (FPGA)” *2011 REDW Record*, [online] Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6062534>
- [15] K. Ellsworth, *Understanding Design Requirements for Building Reliable, Space-Based FPGA MGT Systems Based on Radiation Test Results*, Master’s Thesis, Brigham Young University, June 2012

- [16] K. Ellsworth, A. Harding, C. Ballew, T. Haroldsen, M. Wirthlin, and B. Nelson, "Radiation Testing of FPGA-Based High-Speed Serial Communication," *RADECS Data Workshop*, Sept. 2012.
- [17] R. Monreal and K. Ellsworth, "MGT Data Review: Impact of an Aurora Protocol," Annual XRTC Meeting, March 1, 2012 [online] Available in: <http://www.xilinx.com/member/space/XRTC2012-Proceedings.zip>
- [18] G. Madias, E. Miller, E. Limberg, B. Rasmussen, K. Kohnen, D. Hansen, A. Le, A. Gonzalez, C. Carmichael, and G. Swift, "SEE Characterization of Xilinx Radiation-Hardened, Space-Grade Virtex-5QV," ReSpace/MAPLD, Albuquerque, NM, November 2010.
- [19] G. Allen, G. Madias, E. Miller, and G. Swift, "Recent Single Event Effects Results in Advanced Reconfigurable Field Programmable Gate Arrays," *2011 REDW Record*, [online] Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6062511>
- [20] E. Miller, G. Madias, G. Allen, G. Swift, and C. Carmichael, "Clock Management Tile SEE Test Results," Annual XRTC Meeting, San Jose, CA, March 1, 2012 [online] Available in: <http://www.xilinx.com/member/space/XRTC2012-Proceedings.zip>